Decision-Based Transcription of Jazz Guitar Solos Using a Harmonic Bident Analysis Filter Bank and Spectral Distribution Weighting

Stanislaw Gorlow, *Member, IEEE,* Mathieu Ramona, and François Pachet

arXiv:1611.06505v1 [cs.SD] 20 Nov 2016

***Abstract*—Jazz guitar solos are improvised melody lines played on one instrument on top of a chordal accompaniment (comping). As the improvisation happens spontaneously, a reference score is non-existent, only a lead sheet. There are situations, however, when one would like to have the original melody lines in the form of notated music, see the Real Book. The motivation is either for the purpose of practice and imitation or for musical analysis. In this work, an automatic transcriber for jazz guitar solos is developed. It resorts to a very intuitive representation of tonal music signals: the pitchgram. No instrument-specific modeling is involved, so the transcriber should be applicable to other pitched instruments as well. Neither is there the need to learn any note profiles prior to or during the transcription. Essentially, the proposed transcriber is a decision tree, thus a classifier, with a depth of 3. It has a (very) low computational complexity and can be run on-line. The decision rules can be refined or extended with no or little musical education. The transcriber's performance is evaluated on a set of ten jazz solo excerpts and compared with a state-of-the-art transcription system for the guitar plus PYIN. We achieve an improvement of 34 % w.r.t. the reference system and 19 % w.r.t. PYIN in terms of the F-measure. Another measure of accuracy, the error score, attests that the number of erroneous pitch detections is reduced by more than 50 % w.r.t. the reference system and by 45 % w.r.t. PYIN.**

I. Introduction

AUTOMATIC music transcription, or audio-to-score conversion, is for many the "Holy Grail" of music informatics thanks to the apparent complexity of the task. The problem is even harder in the presence of multiple melodic voices. Over the years, a plethora of different approaches has been proposed; too many to name them all. Some worthwhile ideas can be found in [2]–[7] and in the references given therein. For monophonic signals, it can be said that YIN [8] has established itself as a standard reference. For the more complex polyphonic signals, where multiple notes may be played by different instruments at the same time, any form of a nonnegative matrix factorization (NMF) won the popular vote; the NMF is usually applied to the spectrogram or its square root [9]. A more recent comparison of some of the state-of-the-art techniques was made in [10] on the basis of piano signals. At first glance, the winning algorithm looks a lot like [11], which then again looks like a probabilistic reformulation of the model in [12] including a noise term.

As of today, one of the more popular tasks in music information retrieval constitutes the detection of the dominant melodic voice in a homophonic musical texture [13]. Many times, however, a music signal in which two or more notes are played simultaneously is referred to as *polyphonic*—without further looking into the semantics. Melody detection has its roots in pitch detection of human voice, which actually means the detection of the fundamental frequency of a harmonic oscillator. As mentioned earlier, the most prominent pitch detection algorithm today is YIN. It was originally designed for speech, but it is generic enough to be applied to (monophonic) music as well. In the context of speech, one may also mention [14]. In their work, the authors manage to outperform YIN by a narrow margin with a much more complex and finely tuned overall system. A more recent variant of a YIN-based pitch-tracking system is PYIN [15]. A longer list of algorithms to perform (monophonic) melody extraction from a polyphonic music signal can be found in [13].

Generally, it can be said that neither YIN nor PYIN perform well if not at all on polyphonic signals. NMF-like approaches are non-musical by design, and as such they often generate non-musical artifacts that require manual intervention or any sort of musical post-processing. Probabilistic approaches become computationally intractable if musical constraints in the form of additional priors are incorporated into them; one will have a hard time trying to find the corresponding formulation that must be both: mathematically *and* musically sound. In many cases, the two objectives are opposed and cannot be achieved simultaneously. For these (and many other reasons), the most recent work that directly relates to this paper's topic pursues a bag-of-features approach to reach a result [16].

The greater part of pitch and melody detection algorithms relies on a score function, which is also termed a salience function to underline the perceptual component of pitch and melody, see [17]. One major drawback of the most commonly used score functions is the appearance of phantom pitch at integer multiples of the fundamental, which as a consequence gives rise to octave errors. Another issue with present-day algorithms, which are capable of polyphony, is that they are trained mostly on piano recordings that miss the subtle nuances of a real performance. If you listen to a jazz guitar recording, apart from picked or plucked notes and chords you will notice arpeggios, hammer-ons and hammer-offs, pitch modulations, fretting noise, etc. Put simply, all the elements that render not

This work was financially supported by the European Commission under grant EU-FP7-318770.

S. Gorlow was with Sony Computer Science Laboratory (CSL) in Paris, 6 rue Amyot, 75005 Paris, France. He is now with Dolby Sweden.

M. Ramona and F. Pachet are with Sony CSL Paris.

This paper has supplementary downloadable material. The material includes jazz guitar solo excerpts (licks) and their respective transcriptions. Contact stanislaw.gorlow@gmail.com for further questions about this work.

only a guitar performance expressive and vivid.

The present work was mainly motivated by the fact that an accessible system for guitar performance analysis, which makes use of YIN in combination with a transient detector at its core, did not provide the desired precision and accuracy, in particular when applied to solo improvisations performed on an electric jazz guitar. Another drawback of the system was that it could not be run on-line, i.e. simultaneously with the actual performance. Hence, we decided to review the problem of pitch detection and pitch representation.

In acoustic music, the vast majority of playing techniques can be associated with distinct acoustical phenomena that relate to the physical properties of the instrument [18], [19]. The number of these techniques is finite and can be reduced further if the context such as the artist or the genre is known. Hence, on condition that we possess a suitable signal representation to measure the concomitant phenomena, we can define a set of rules to identify the techniques employed by the performer and decipher the musical message. This is one of the central ideas behind the present paper.

The spectrogram has emerged as the standard representation for speech (also called a *voicegram*), music, and other signals that exhibit a structure along the frequency axis, which may also change over time [20]. However, for tonal music in particular, the key question is often not about the physical composition of the signal in terms of frequency components, but rather about its musical composition in terms of pitch or melody. And so, this work is not so much about pitch or melody detection as it is about finding a better representation for the harmonic content of a music signal, which as a natural consequence would facilitate the analysis and the transcription task.

The major contribution of the present work consists in the proposed filter bank. The developed time-pitch transform may be viewed as a multi-pitch score function that varies over time. This type of representation we call a *pitchgram*. A genuine pitchgram can assist a trained musician during transcription tasks or it can help an amateur find the melody, the solo, or the progression of chords in the accompaniment. A more reliable score function should also improve the accuracy of the salience-based detection algorithms mentioned in [13]. In addition, we make public the used data set with the manual transcriptions for everybody interested in this area of research.

To substantiate our statements in the preceding paragraphs, we manually construct a decision tree out of empirical decision rules that translate into musical acoustics. The tree serves as a pitch classifier, i.e. an automatic transcriber, while resorting to any of the two variants of the newly developed time-pitch representation for tonal signals as input. We demonstrate that our new time-pitch representation is more reliable than, e.g., the (note) activation diagram of an NMF. On a final note, we would like to mention some similarities and differences w.r.t. the so-called "specmurt" analysis [21]. Both methods do not require prior knowledge of the number of latent components, i.e. distinct pitches. Both are sensitive to signal amplitude: in the presence of strong components, weak components vanish due to normalization and thresholding. The advantages of the filter bank approach with distribution weighting are:

- The pitch axis can be neither linear nor logarithmic;
- Each harmonic series is weighted differently, i.e. there is no need to compute a common harmonic weighting;
- There are no instability issues related to inverse filtering or deconvolution and there are no iterations;
- The computation of the pitchgram is less costly, since it requires only a single (forward) Fourier transform in the case of the faster variant.

The structure of the paper is as follows. In Section II, the definition of pitch on a twelve-tone equal temperament scale is briefly reviewed in order to formulate the properties of the corresponding frequency spectrum. Inspired by the physics, a multi-tone pitch analysis filter bank is designed in Section III. The output from the filter bank is a time-pitch representation of the signal: the pitchgram. The latter is used as input for a rule-based pitch classifier, which is detailed in Section IV. In the following Section V, the pitch classifier is then evaluated and compared. Section VI concludes the paper.

II. Pitch on a Twelve-Tone Equal Temperament Scale

A. Pitch and Harmonic Spectrum

Pitched musical instruments produce sounds, or notes, of definite pitch that can be discerned by a listener. Sounds with definite pitch usually possess a harmonic spectrum with many or only a few components. This is because a musical tone that is generated on a pitched instrument has several modes of vibration that occur simultaneously. In steady state, a pitched instrument can hence be approximated as a harmonic oscillator, such as a string, which vibrates at the frequencies of a harmonic series.[1] The listener hears all frequencies at once. The frequency of the slowest vibration is the fundamental frequency, which is usually perceived as the musical pitch. The other frequencies are overtones. A harmonic series consists of frequencies that are integer multiples of the fundamental frequency, which is equally the first harmonic. The overtones start at the second harmonic. Harmonics are also called partials, referring to the different parts of the harmonic spectrum of a musical tone. Harmonics have the property that they are periodic at the fundamental frequency, and so is their sum. In what follows, the pitch is quantified as the fundamental frequency of a harmonic series, even though it is a subjective psychoacoustical attribute rather than a purely objective physical property [23].

B. Equal-Tempered Chromatic Scale

The chromatic scale is a musical scale with twelve tones or pitches [24]. On an *equal-tempered* instrument, all the semitones, i.e. the intervals between two adjacent tones, have the same size. As a consequence, the notes of an equal-tempered chromatic scale are equally spaced. More precisely, the twelve-tone equal temperament divides the octave, which is the interval between two consecutive harmonics, into twelve

[1] Some musical instruments exhibit inharmonicity. However, inharmonicity is negligible for the nylon strings of guitars. It only becomes appreciable for thick steel strings, as in the piano [22].

segments, all of which are equal on a logarithmic scale. Equal-temperament tuning systems are widely used in contemporary music. The MIDI tuning standard, e.g., by default expresses pitch on that scale.[2]

III. Pitch-Synchronous Chromatic Bident Analysis Filter Bank

In this section, a novel pitch-analysis filter bank is designed. It has for purpose the decomposition of a (tonal) signal into a set of pitch coefficients, each one being associated with a distinct tone or note. In principle, the signal's texture can be monophonic, polyphonic or homophonic, albeit the latter case is treated only marginally. The filter bank is shown in Fig. 1. Its constituent blocks are discussed in the following sections. The particularity of the proposed filter bank lies in its ability to alleviate octave errors on the hypothesis that a tone spectrum is harmonic and has a decaying tendency, which can facilitate further analysis.

A. Pitch-Synchronous Comb Filter and Harmonicity Coefficient

1) Pitch-Synchronous Comb Filter: Consider the relation between a tonal input signal $x(t)$ and the corresponding output signal $\tilde{x}(t)$ to be defined as [25]

$$\tilde{x}(t) = x(t) + a\,x(t - T_0), \tag{1}$$

where T_0 is the pitch period and $a \in [-1, 1]$ is a scaling factor. The difference equation (1) describes the *feed-forward* comb filter. Its frequency response consists of a series of regularly spaces notches, hence the name. It is based on the principle of constructive and destructive interference between two waves that are correlated or coherent with each other. The alternative *feed-backward* comb filter is described by

$$\tilde{x}(t) = x(t) + a\,\tilde{x}(t - T_0). \tag{2}$$

Note that the feed-backward comb filter variant is stable only if $|a| < 1$. The magnitude response of the filter for a delay equal to the pitch period of A_3 is shown in Fig. 2. It can be seen that the spikes appear at integer multiples of the fundamental including the zero frequency (direct current). Each channel of the analysis filter bank has an optional comb filter at the entrance that resonates at the nominal frequency and the corresponding integer multiples. Accordingly, the nominal frequency in Fig. 2 is the frequency of A_3, which is 220 Hz. The purpose of the pre-filtering is to accentuate the components of a complex tonal mixture at the frequencies of a particular harmonic series and to attenuate all other components. In the remaining part of the manuscript, we resort to the feed-backward variant of the comb filter.

2) Harmonicity Coefficient: To assess the presence of a particular harmonic series in the input signal $x(t)$, we introduce a *harmonicity coefficient*. We define it as the root mean square (RMS) of (desirably) one period of the comb-filtered signal $\tilde{x}(t)$,

$$\eta(T_0) \triangleq \sqrt{\frac{1}{T_0}\int_0^{T_0} [\tilde{x}(t)]^2\,\mathrm{d}t}. \tag{3}$$

[2]See http://www.midi.org/

The output of the first block is the original input signal $x(t)$ weighted by its harmonicity coefficient w.r.t. the nominal frequency of the respective filter bank channel, not the filtered signal:

$$y(t, T_0) = \eta(T_0)\,x(t). \tag{4}$$

B. Autocorrelation Function and Compression

1) Autocorrelation Function: The second block initially computes the autocorrelation function (ACF) of the pre-weighted signal $y(t)$, which consists of multiple decaying complex tones with distinct harmonic spectra. The ACF is typically defined as the limit of the time average,

$$r_{yy}(\tau) = \lim_{T\to\infty} \frac{1}{T}\int_0^T y(t+\tau)\,y(t)\,\mathrm{d}t, \tag{5}$$

where τ denotes the time lag and T is the observation window, respectively. The mean signal power is given by the ACF at zero lag, i.e.

$$r_{yy}(0) = \lim_{T\to\infty} \frac{1}{T}\int_0^T [y(t)]^2\,\mathrm{d}t. \tag{6}$$

When normalized by the mean signal power, the autocorrelation function is often referred to as the autocorrelation coefficient (ACC)

$$\rho_{yy}(\tau) = \frac{r_{yy}(\tau)}{r_{yy}(0)}. \tag{7}$$

Should the ACC and not the ACF be computed in the second block, the harmonic weighting can be skipped. This is why (4) is optional. Two important properties of the ACF (and the ACC) are that it is (a) real-valued and (b) evenly symmetric. Hence, its Fourier transform is real-valued and evenly symmetric as well. These properties are exploited in the next section, where the chromatic bident filter is derived.

2) Compression: In quite a few papers on polyphonic music transcription, see e.g. [26], it was observed that often better results are obtained if the magnitude spectrum is used instead of the power spectrum along the frequency dimension. And since the power spectrum is the Fourier transform of the ACF, we achieve a similar effect by taking the square root of $\left|r_{yy}(\tau)\right|$ while keeping the sign, i.e.

$$\tilde{r}_{yy}(\tau) \triangleq \mathrm{sgn}[r_{yy}(\tau)]\,\left|r_{yy}(\tau)\right|^{1/2}, \tag{8}$$

where sgn denotes the signum or sign function. We call this *compression*.

C. Cosine-Modulated Chromatic Bident Filter

1) Cosine-Modulated Chromatic Sinc Filter: Imagine that we wanted to design a filter that has its peak or maximum where is the fundamental frequency of one particular tone and zeros at the fundamental frequencies of all other tones in a given range. An ideal filter that fits the description is a shifted Dirac delta function. However, since its impulse response, which is equivalent to the size of the observation window, is infinite, the Dirac delta filter is impracticable. If we truncate the window, the Dirac delta becomes the sinus cardinalis or sinc for short. The width of the main lobe of the sinc function is inversely proportional to the window size, while the zero

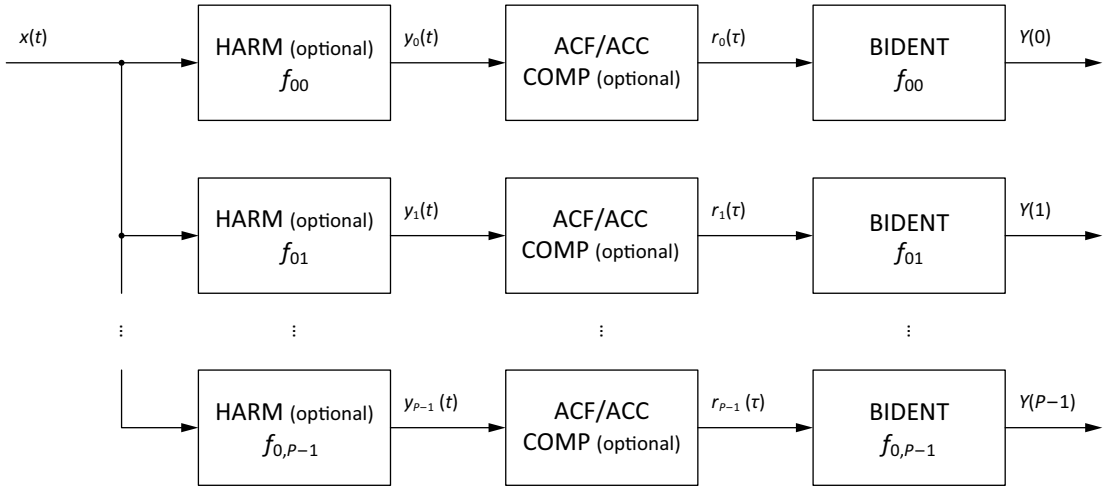

Fig. 1. P-tone pitch-synchronous chromatic bident analysis filter bank.

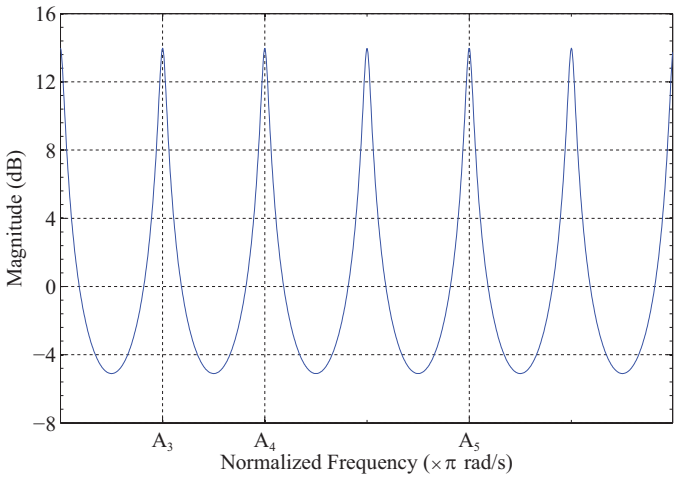

Fig. 2. Feed-backward comb filter response for A_3 ($a = 0.8$).

crossings of the normalized sinc occur at non-zero integers; the zero crossings are *equidistant*. Therefore, if we assume that each tone in a tonal mixture carries most of its energy in the harmonic partials, we can dimension a rectangular window in such a manner that the first zero crossing is *one twelfth* of an octave away from the fundamental frequency of the tone of interest in the lower octave, i.e. between $f_0/2$ and f_0; one twelfth because we are considering a chromatic scale. This is to prevent lower-pitched and also higher-pitched tones from leaking significantly into the filtered signal, which ideally would oscillate at its fundamental frequency only. Note that the number of zero crossing in the upper octave, i.e. between f_0 and $2\,f_0$, is *twice* as large. Attention should also be paid to the fact that the leakage from neighboring tones is not perfectly suppressed. This is because the zero crossings of the sinc filter are equidistant on a linear scale, whereas the pitches are equidistant on a logarithmic scale. The effect is negligible, however, as the side lobes of the sinc decay rather quickly. This is further supported by the fact that most of the energy is concentrated under the main lobe with the area under the side lobes adding up to a small (negative) value.

The window discussed above is formally given by the rectangle function, which is defined as

$$\mathrm{rect}\left(\frac{t}{24\,T_0}\right) = \begin{cases} 1 & \text{if } |t| \leqslant 12\,T_0, \\ 0 & \text{otherwise,} \end{cases} \tag{9}$$

where T_0 is the pitch period of the tone of interest. Accordingly, its (unitary) Fourier transform is given by [27]:

$$\mathrm{rect}\left(\frac{t}{24\,T_0}\right) \circ\!\!-\!\!- \; 24\,T_0\,\mathrm{sinc}\left(24\,\frac{f}{f_0}\right) \tag{10}$$

with

$$\mathrm{sinc}\left(24\,\frac{f}{f_0}\right) = \frac{\sin(24\,\pi\,{}^{f}\!/_{f_0})}{24\,\pi\,{}^{f}\!/_{f_0}}, \tag{11}$$

where $\circ\!\!-\!\!-$ symbolizes the direction of the mapping: from a *discrete* (time) domain $\circ$ to a *continuous* (frequency) domain. From (11) we can immediately see that the sinc has twelve zero crossings in the lower octave and twenty-four of them in the upper octave, respectively. In the previous section it is underlined that the Fourier transform of the ACF is real. Now, if we window (5) by (9) and then apply the Fourier transform, we obtain a normalized cosine transform at f_0, which again is real,

$$\tilde{Y}(f_0) \triangleq \frac{1}{24T_0}\int_{-\infty}^{\infty} \mathrm{rect}\left(\frac{\tau}{24T_0}\right) r_{yy}(\tau)\cos(2\pi\,f_0\,\tau)\,\mathrm{d}\tau. \tag{12}$$

The normalization term before the integral makes sure that the integral is *independent* of the window size, and so the pitch period T_0. Note that (12) is tantamount to a *zero-lag cross-correlation* between the cosine-modulated rectangle window and the ACF of the comb-filtered input signal $x(t)$. But what is more: because the rectangle, the ACF, and the cosine are all even functions, the cross-correlation between the above terms is identical with their *convolution*. Therefore, (12) can be interpreted as a *filtering operation* on the Fourier transform of the ACF, i.e. the harmonically weighted input power spectrum, by a pitch-shifted sinc. In the frequency domain, this corresponds to a multiplication of the signal

spectrum by the sinc function. Finally, due to symmetry, (12) simplifies to

$$\tilde{Y}(f_0) = \frac{1}{12T_0}\int_0^{12T_0} r_{yy}(\tau)\cos(2\pi f_0\,\tau)\,\mathrm{d}\tau. \tag{13}$$

2) Cosine-Modulated Bident Window: In the previous paragraph it was shown how a particular partial from a harmonic series can be filtered out, while suppressing simultaneously the neighboring tones. But is there a means to demarcate an overtone from the fundamental frequency of the same harmonic series? Or in other words, how can we *score high* the *first partial*, while *scoring low* all the *other partials* at the same time; and to make it more difficult, by the use of a simple filter?

Consider the following prototype function:

$$g(t) = \alpha\sin(3\pi f_0 t)\tan(\pi f_0 t) - \beta. \tag{14}$$

If we modulate the zero-frequency carrier in (14) by half the fundamental frequency of a tone, we obtain

$$h(t) \triangleq g(t)\cos(\pi f_0 t), \tag{15}$$

which is the periodic (infinite) impulse response of our filter. Its frequency response is given by the Fourier transform, which is

$$H(f) = \tfrac{\alpha}{4}\left[\delta(f-f_0) - \delta(f-2f_0)\right] - \tfrac{\beta}{2}\,\delta(f-f_0/2), \tag{16}$$

consisting of three weighted Dirac deltas with different signs at the frequencies $f_0/2$, f_0, and $2f_0$. If we multiply (15) by (9) to limit the length of the filter, it becomes

$$\begin{aligned}\alpha\,6T_0\left[\operatorname{sinc}\left(24\,\frac{f-f_0}{f_0}\right) - \operatorname{sinc}\left(24\,\frac{f-2f_0}{f_0}\right)\right] \\ -\beta\,12T_0\operatorname{sinc}\left(24\,\frac{f-f_0/2}{f_0}\right) \multimap \operatorname{rect}\left(\frac{t}{24T_0}\right) h(t),\end{aligned} \tag{17}$$

where $\multimap$ now symbolizes the inverse mapping. In the case where $\alpha = 2$ and $\beta = 1$, the three spikes have the exact same magnitude, see Fig. 3. Due to the shape of the truncated window $w(t)$ in the frequency domain,

$$w(t) \triangleq \operatorname{rect}\left(\frac{t}{24T_0}\right) g(t), \tag{18}$$

it is given the name *harmonic bident*. From Fig. 3(b) we see that between two adjacent spikes there are eleven zero-crossings in the lower octave and twenty-three in the upper octave. Accordingly, with the new window, (12) becomes

$$Y(f_0) \triangleq \frac{1}{24T_0}\int_{-\infty}^{\infty} w(\tau)\,r_{yy}(\tau)\cos(\pi f_0\,\tau)\,\mathrm{d}\tau, \tag{19}$$

or more compactly written

$$Y(f_0) = \frac{1}{12T_0}\int_0^{12T_0} r_{yy}(\tau)\,h(\tau)\,\mathrm{d}\tau \tag{20}$$

with $T_0 = 1/f_0$.

The rationale behind the bident filter is the following. A pitched instrument produces a harmonic spectrum that is assumed to be (exponentially) *decaying*. Broadly speaking, this is motivated by the second law of thermodynamics. So, if the holder of the bident, i.e. the middle spike, is shifted to the pitch of a complex tone, the output of the filter contains all the energy of the fundamental, which is reduced by the energy of the first overtone (right spike). If the spectrum has a decaying trend, the output is positive. However, if the holder is placed on any of the overtones, the filter output should be negative, because the energy of the overtone is now reduced not only the the *next* overtone but also by the *previous* overtone (left spike), which in the case of the first overtone is the fundamental. Due to the falling trend of the harmonics, one may anticipate that most of the overtones will have a *negative* output and that the pitch will be the frequency with the largest *positive* correlation or *score*. It should be noted, however, that the spectrum (of overtones) may not always be strictly monotonically decreasing. This can depend on the timbre of the instrument and the playing technique, which can put the accent on the first or the first few overtones rather than the fundamental. In that case, the fundamental might have a lower (positive) score than the accentuated overtones. Then, it is up to the transcriber to select the lowest partial of a harmonic series with a significant score. Any constructive interference effects between multiple harmonic series (polyphony) are not

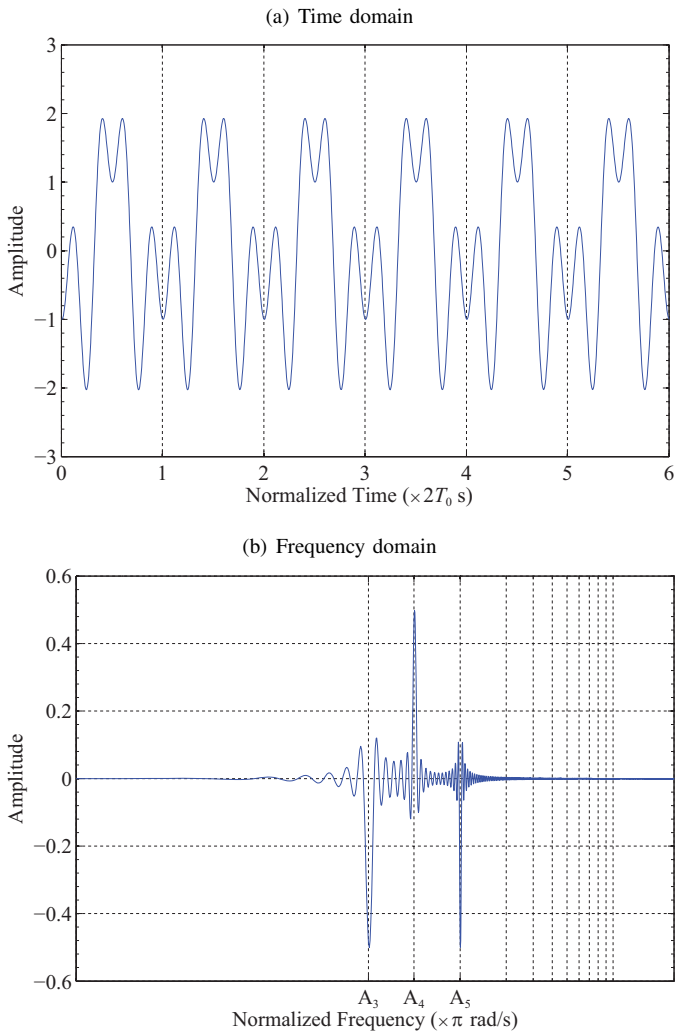

Fig. 3. (a) Cosine-modulated bident window and (b) the corresponding filter response for A_4.

explained by this model. In [28], e.g., the authors promote a somewhat similar idea: a comb with alternate polarity, which should minimize octave errors.

D. Pitch Analysis of Non-Stationary Signals

The extension of (19) or (20) to a succession of tones is very much straightforward. You just choose the resolution of the time grid and apply (20) at all time instants for all possible values of f_0. Formally, this is stated as

$$Y(m,f_0)=\frac{1}{24T_0}\int_{-\infty}^{\infty}w(\tau)\,r_{yy}(m,\tau)\cos(\pi f_0\,\tau)\,\mathrm{d}\tau, \tag{21}$$

or equivalently

$$Y(m,f_0)=\frac{1}{12T_0}\int_{0}^{12T_0}r_{yy}(m,\tau)\,h(\tau)\,\mathrm{d}\tau, \tag{22}$$

where $m\equiv mT_g$ (equivalent to sampling), $m\in\mathbb{N}_0$, and

$$r_{yy}(m,\tau)=\lim_{T\to\infty}\frac{1}{T}\int_0^T y(mT_g+t+\tau)\,y(mT_g+t)\,\mathrm{d}t. \tag{23}$$

T_g denotes the resolution of the time grid. The resulting time-pitch diagram is what we call the *pitchgram*. T_g can be chosen as, e.g., the duration of a thirty-second note. It can be noted that a corresponding chromagram [29] is easily obtained from (22) by

$$Z(m,\zeta)=\sum\nolimits_{\mu\in\mathbb{N}^+}Y\big(m,\mu\cdot f_{0,\zeta}\big), \tag{24}$$

where ζ denotes a *pitch class* or chroma and μ is the *octave number* or tone height. It should be of greater benefit to use the rectangle window, i.e. the sinc filter, in that case. Also note that $f_0\in\mathbb{R}$ by definition, i.e. it may take on any value and is not restricted to any tuning standard.

E. Discrete-Time Discrete-Pitch Implementation

For the realization of (22) as a computer program, the time variable t or τ must be discretized. But first, for the sake of convenience, the frequency variable f_0 is substituted for the MIDI note number p, i.e.

$$Y(p)=Y\left(440\cdot 2^{p-69/12}\right) \tag{25}$$

for $p=0,1,\ldots,127$ and f_0 in Hz. The MIDI note with the number $p=69$ belongs to the concert A ($f_0=440$). The middle C lies nine semitones below ($p=60$). Finally, the values of $Y(p)$ are evaluated at $t=n\equiv nT_s$, $n\in\mathbb{Z}$, where T_s denotes the sampling period. The pitch period of a note T_0 is substituted for $N_0\equiv N_0T_s$, respectively. It should be noted that the pitch mapping in (25) is arbitrary and not binding. In other words, the pitch resolution can be finer than a semitone with any tuning other than 440 Hz for A and the pitch range can be limited to whatever makes sense w.r.t. the instrument.

With the above definitions, the feed-backward discrete-time comb filter reads

$$\tilde{x}(n)=x(n)+a\,\tilde{x}(n-N_0), \tag{26}$$

and the harmonicity coefficient computes as

$$\eta=\sqrt{\frac{1}{N_0}\sum_{n=0}^{N_0-1}[\tilde{x}(n)]^2}. \tag{27}$$

Respectively, the discrete autocorrelation is defined as

$$r_{yy}(\nu)=\lim_{N\to\infty}\frac{1}{N}\sum_{n=0}^{N-1}y(n)\,y(n-\nu) \tag{28}$$

for $\nu=0,1,\ldots,12N_0-1$. The cosine-modulated bident becomes

$$h(n)=g(n)\cos\left(\frac{\pi n}{N_0}\right) \tag{29}$$

with

$$g(n)=\alpha\sin\left(\frac{3\pi n}{N_0}\right)\tan\left(\frac{\pi n}{N_0}\right)-\beta \tag{30}$$

for $n=0,1,\ldots,12N_0-1$. Taking everything into account, the time-dependent score for the mth tone is given by

$$Y(m,p)=\frac{1}{12N_0}\sum_{n=0}^{12N_0-1}r_{yy}(m,n)\,h(n). \tag{31}$$

The corresponding MATLAB code is part of the downloadable material.

F. Example

Fig. 4 shows the results from the pitch analysis of a jazz guitar solo. In Fig. 4(a) we see the signal's log-spectrogram. The ideal binary mask[3] is shown in Fig. 4(b). Figs. 4(c)–4(f) illustrate the different variants of the pitchgram. The left column shows the results obtained with the sinc filter, while the right column shows the results obtained with the bident filter, respectively. In the middle row, both the pitchgrams were computed based on a harmonically weighted ACF. In the lower row, the ACF is normalized by the mean signal power. Furthermore, the ACF was compressed in all four cases. The spectrogram was computed using a 4096-point discrete Fourier transform and a Hamming window of equal length with an overlap of 75 % between consecutive blocks. The pitchgram uses the same time resolution.

Looking at the pitchgram, one can easily identify single notes, arpeggios, and chords. At various time instants, the pitch is modulating between the unison and the minor second (one semitone)—a clear indication of vibrato, see Fig. 4(a). For the most part, this translates to amplitude variations in the pitchgram. The issue with vibratos is how not to interpret (periodic) variations in amplitude as a tremolo, which would otherwise make a single note appear as several notes in the transcription. As a similar case, hammer-ons, which may be not strong enough to be recognized as note onsets, and pull-offs, which may not make the tone decay to a low power level before the string is struck again, both require special attention. This brings up the question whether, to a certain degree, the affected notes could be associated with a playing technique, such as vibrato, legato, etc. If this was the case, one could then carry out a more sophisticated performance analysis with regards to articulation and expression, something that is beyond pitch and onset. One may also try to identify technical errors, which otherwise would result in phantom notes. The latter are provoked by fast finger moving and sloppy finger

[3]The binary mask was obtained by fine tuning an automatic transcription by hand and ear in a digital audio workstation.

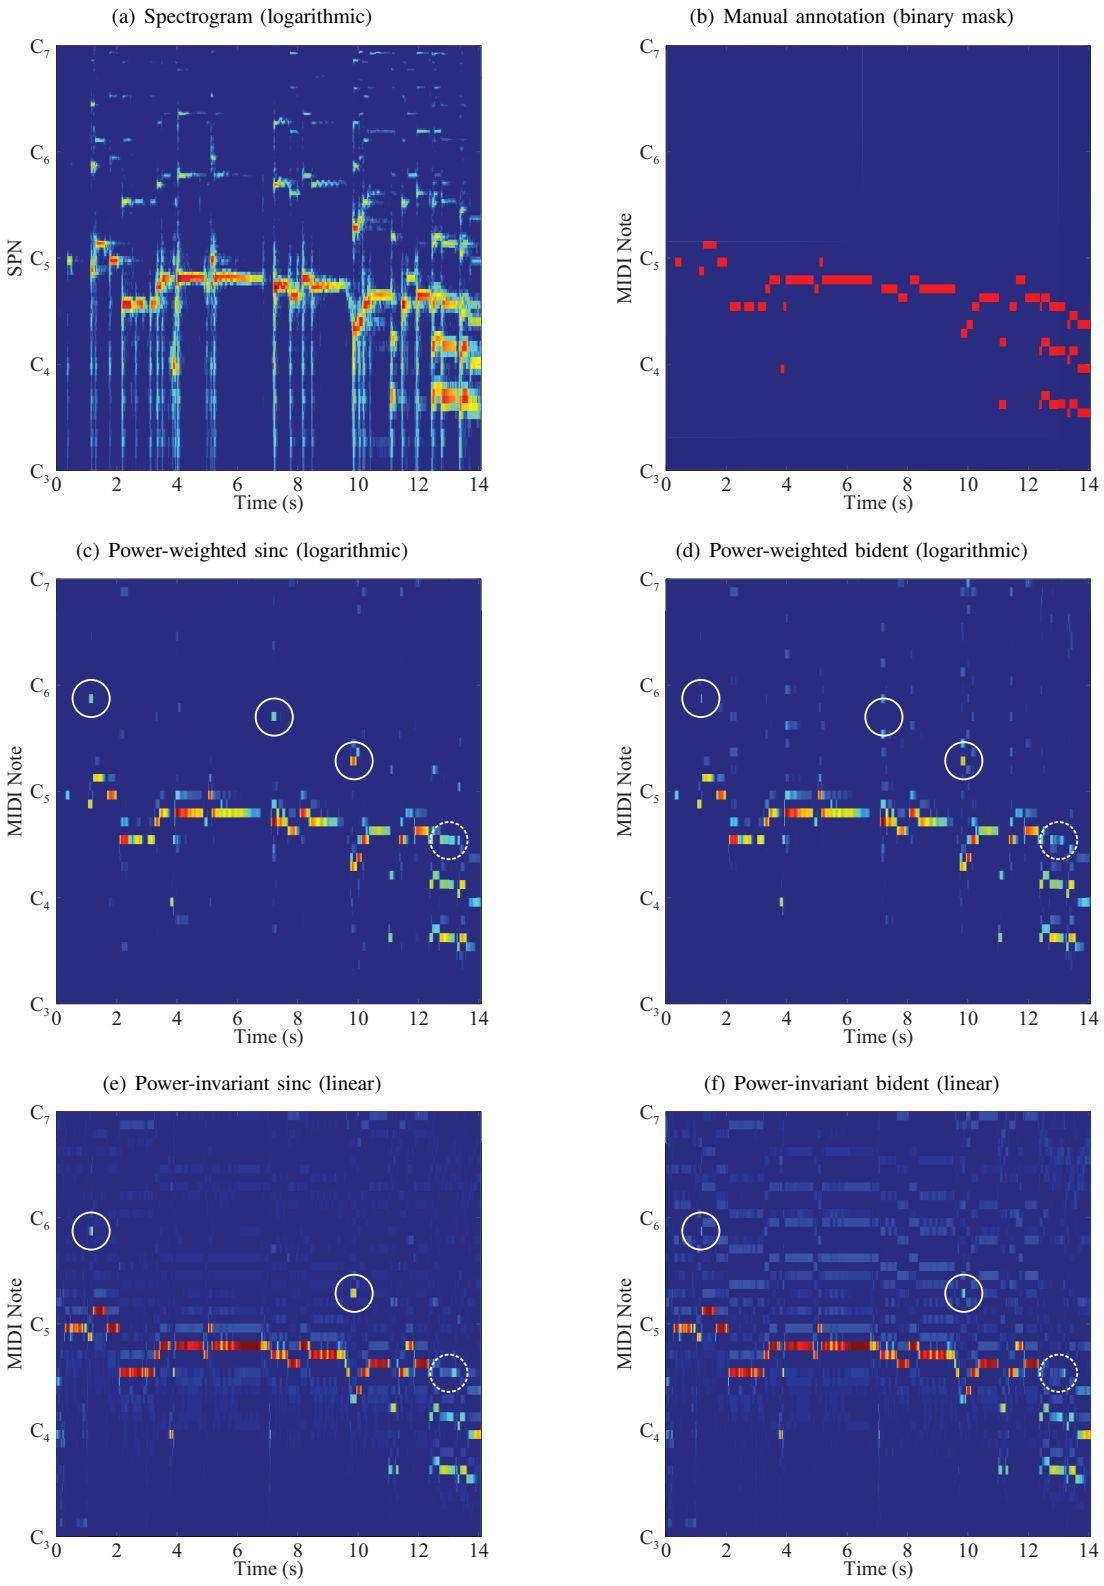

Fig. 4. Results from the pitch analysis of a jazz guitar solo ("Solar").

placement when plucking. It can also be seen that sliding chords (glissando) are more tricky to handle, because the note onsets are slurred and the fingers encroach on neighboring tones.

The bident successfully eliminates, or at least alleviates, the octave ambiguities that are due to an accentuated (first) overtone, see white circles. Occasionally, this also happens at the expense of the fundamental, which might perish together with

the overtones; the dashed circle points to such a case where the upper forth disappears almost completely. In part, this type of effect is due to the chosen signal representation: the ACF. An overtone can come out stronger (in the spectrogram) than it is actually heard, when the phase relations between the harmonics are impure. Certainly, we can also discover some artifacts, which are a consequence of the side lobes of the sinc filter and the resulting ripple of the bident. As a final remark, we would like to mention that note velocity is useful for the assessment of dynamics in a performance or for the detection of a particular form of accentuation.

G. Comparison with NMF

To quantify the improvement that one can achieve with the filter bank approach in comparison to an established tool for (polyphonic) music transcription, such as the NMF, we carry out the following experiment. Firstly, we compute the power-weighted and the power-invariant pitchgram, as in Figs. 4(d) and 4(f), for the corpus we resort to later in our performance evaluation, see Section V-B. The two representations are then multiplied by each other and the result is normalized by the Frobenius norm. In a way, this is equivalent to computing the posterior pitchgram (the power-invariant pitchgram may also be interpreted as the tonality likelihood). The new pitchgram is then smoothed, see Section IV-A, and finally compared to a threshold with a binary decision, to yield a binary mask. It should be noted that the power-weighted pitchgram accounts for the duration of notes. Secondly, the NMF is computed. A standard procedure is to take the magnitude spectrum for the time-frequency representation (TFR) and to use the Kullback–Leibler divergence as distance metric for the factorization. In this regard, we use Lee and Seung's (original) algorithm with multiplicative update rules from the NMFlib.[4] The algorithm is run ten times (with different initializations) for each signal and the factorization with the smallest error is kept. The rank of the factorization is set equal to the number of notes in the signal, which is obtained from the reference transcription. In an intermediate step, the activation matrix is converted into a time-pitch representation and then post-processed in the exact same manner as the pitchgram. The frequency resolution for the NMF is 10.8 Hz per bin and the time resolution for both the pitchgram and the NMF is 1024 samples (23.2 ms). The outcome of the comparison is shown in Fig. 5. On each box, the central mark is the median, and the edges of the box are the 25th and 75th percentiles. The whiskers indicate the most extreme data points, while outliers are plotted as circles. The F-measure is defined in Section V-C.

Quite clearly, the posterior pitchgram is way more reliable than the NMF under the employed protocol. The result can be explained as follows. Firstly, the error of the NMF is defined as the (statistical) distance between the original spectrogram and its factorization—not on the factors. Given that the rank, i.e. the number of factors, was set equal to the ideal number of perceptually distinct notes, pitch modulations within half a semitone can easily result in wrong pitch. This is because the spectral note profile of a neighboring tone can be closer to a pitch-bended version of the original tone, or simply because the spectral imprint of a similar note gives a smaller error in reference to the original spectrogram. Secondly, it is common knowledge that the attack (transient) and sustain (tonal) parts of a note should be represented with different bases. If only the tonal part can be assigned to a basis, note onsets become another source of error. Finally, it should be noted that better results can be obtained, if the NMF bases are pre-learned on isolated note spectrograms with different accentuations and/or playing techniques. The advantages of the pitchgram are:

- Number of notes (rank) not needed as prior knowledge;
- No ambiguity in note profiles (dictionary);
- No convergence issues (convexity/initialization);
- No distinction between attack and sustain necessary.

[4] http://www.ee.columbia.edu/~grindlay/code.html

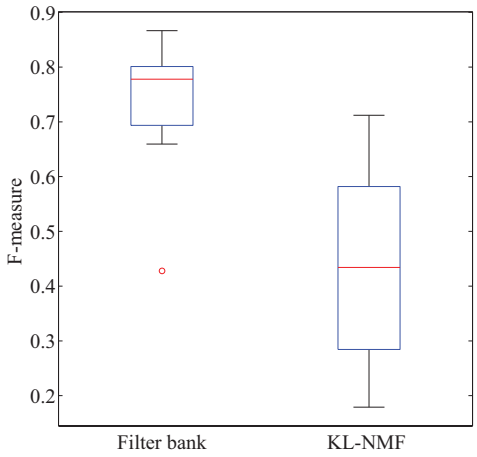

Fig. 5. Pitch detection accuracy: comparison between the filter bank and the KL-NMF with multiplicative updates w.r.t. the F-measure.

IV. Automatic Transcription

In this section, we provide some technical details on how to realize an *exemplary* transcriber that resorts to the pitchgram. To counter the issue of signal level dependency of the tuning parameters, we suggest to normalize the input signal level to a reference signal level and to perform the detection steps on both the power-invariant and the power-weighted pitchgrams, whenever ambiguities arise. The complexity of the transcriber and its robustness is largely determined by the onset and decay detection blocks.

A. Pitch Detection

The detection of a note should be reliable, and so it requires a robust decision criterion. With this goal in mind, we define the following normalized derivative:

$$\dot{Y}(m,p) = \frac{\Delta Y(m,p)}{\overline{Y}(m,p)} = \frac{Y(m,p) - Y(m-1,p)}{[G(\cdot) * Y(\cdot,p)](m)}, \qquad (32)$$

where $\Delta Y(m,p)$ is the first-order derivative of Y and $\overline{Y}(m,p)$ is the smoothed value of Y in (m,p). G is the smoother and $*$ denotes linear convolution, respectively. One can also use the median filter in place of the linear smoother. A note onset is subject to the following conditional (if-then) statement:

$$Y(m,p) > \varepsilon_1 \wedge \dot{Y}(m,p) > \varepsilon_2 \rightarrow \text{onset}(m,p) \qquad (33)$$

with ε_1 and ε_2 being empirical threshold values. So, an onset is defined by a sufficiently high pitch score and a steep slope. In addition, one may also check whether the slope is rising.

1) Transient Noise: In most cases, the onset of a note is accompanied by transient noise, which spreads over the neighboring pitches. This is especially the case for the pitchgram that was obtained from the spectrogram. To make sure that transient noise is not mistaken for a note onset, we check whether the following condition is fulfilled:

$$\Delta Y(m,p) > \Delta Y(m,p-1) \wedge \Delta Y(m,p) > \Delta Y(m,p+1) \quad (34)$$

subject to the existence of $p-1$ and $p+1$. By transient noise we mean the score spread that appears at the onset of plucked or picked notes.

2) Fretting Noise: Another form of noise that can falsify the detection of a note onset is fretting noise. Fretting noise appears when the finger is moved along the guitar neck and the string is pressed down behind a fret. It is also typical for hammer-ons and pull-offs. As the origins of fretting noise are different from the origins of transient noise, we introduce another condition to make sure that fretting noise is not mistaken for pitch, which is

$$Y(m,p) > Y(m,p-1) \wedge Y(m,p) > Y(m,p+1) \quad (35)$$

subject to the existence of $p-1$ and $p+1$. Just as (34), (35) implicitly assumes that the tonality of noise is lower than the tonality of the actual pitch.

3) Octave Errors: A simple method to avoid octave errors in monophonic music material is to ensure that there was no note already detected twelve semitones or one octave below the current pitch value p. In that case, the new onset should be discarded.

4) Note Velocity: According to the MIDI standard, velocity is a measure of how rapidly and forcefully a key on a keyboard is pressed when it is idle. In the same way, we may describe with which force a string was hit at note onset:

$$\text{onset}(m,p) \rightarrow \text{velocity}(m,p) \propto Y(m,p), \quad (36)$$

where $Y(m,p)$ is the power-weighted pitchgram from (31). In (36), it is assumed that the velocity (force) is proportional to the measured power at onset. Another way to obtain it, is by integrating the energy over the entire note duration.

B. Pitch Tracking

After detection of an onset, the respective note is deemed to be active, i.e. key pressed or string vibrating, until its energy falls below a certain threshold value and the first derivative is sufficiently large and negative, see (32). And so, the condition for a note decay is as

$$Y(m,p) < \varepsilon_3 \wedge \dot{Y}(m,p) < \varepsilon_4 \rightarrow \text{decay}(m,p), \quad (37)$$

where ε_3 and ε_4 are again empirical threshold values. It may further be wise to check whether the slope in falling to avoid ambiguities. Accordingly, the note duration is given by

$$d(p) = m_{\text{off}} - m_{\text{on}}, \quad (38)$$

where m_{on} is the onset instant and m_{off} the decay instant.

1) Vibrato: On various occasions, the player may make use of vibrato as a means of musical expression. Depending on the hop size of the sliding DFT, i.e. the time resolution of the pitchgram, and the rate of vibrato, the pitch may appear and disappear with a certain regularity. Given (37) only, vibrato will be transcribed almost surely as tremolo, i.e. a quick repetition of the same note. Since the typical extent of vibrato is less than a semitone either side of the note, we should avoid mistaking vibrato for tremolo by replacing the first hypothesis in (37) by

$$Y(m,p) + Y_{\text{vib}}(m,p) < \varepsilon_3, \quad (39)$$

where

$$Y_{\text{vib}}(m,p) = \max\{Y(m,p-1), Y(m,p+1)\} \quad (40)$$

on the supposition that there are no active notes one semitone above and below the current pitch value p.

2) Pruning: Once we made the decision that a note is no longer active, we might verify that is not too short, and hence not an artifact of an imperfect time-pitch representation. We call it pruning in reference to machine learning terminology. The condition for a note to be valid is simply

$$d(p) > d_{\text{min}}. \quad (41)$$

It should be noted that if d_{min} is too long, notes belonging to an arpeggiated chord might get lost as well.

C. Transient Detection

The condition in (33) for the note onset may not always be sufficient. This is usually the case when one and the same note is rapidly reiterated (tremolo) or, in more general terms, when the string is picked or plucked again before it stopped vibrating at a high amplitude. In such a case, the slope of Y in point (m,p) is not steep enough for the second hypothesis in the conditional statement to be true. Then, it is helpful to compute a transient score by marginalizing out the pitch from $Y(m,p)$,

$$Y(m) = \sum\nolimits_p Y(m,p) \qquad \forall\, p, \quad (42)$$

and by evaluating the normalized derivative of the marginal score $\dot{Y}(m)$ in addition to (33). It must be admitted, however, that an additional transient detector might still be necessary.

V. Performance Evaluation

In this section, we present the results of the performance evaluation of our rule-based transcriber in regard to another state-of-the-art transcriber for the guitar. Given the complexity and the number of processing steps, or layers, of most of the melody detection algorithms, we would consider it rather unfair to make a comparison between those systems and our transcriber, which operates *exclusively* on the pitchgram. In addition, we would like to compare our transcriber to a system that was *tuned to the guitar*, and not the piano, and which does not require prior training. We also include PYIN.

A. Algorithms

In our evaluation, we compare the output of the following transcription systems:

A The reference transcriber [30], [31];
B Our (decision-based) transcriber with the pitchgram being computed using the fast Fourier transform;
C The same transcriber as in *B*, but with the original time-domain pitchgram as input (see Section III);
D PYIN [15], which is an improved version of YIN.

The reference transcriber is an updated version of the extraction module used in the two references. It was tested with both *nylon and steel* guitar strings. The extraction module combines the outputs of an onset detector and a pitch detector. The pitch detector is YIN. The pitch trajectory in between two note onsets is mapped to a single pitch value in order to obtain a piano-roll-like representation of the input signal. The onset detector is described in [32]. It should be noted that the algorithm is of great value for the comparison, because 1) YIN is the de-facto standard in most monophonic transcription tasks and 2) given the low degree of polyphony of our test material, it is expected to perform well. The same can be said about the onset detector. PYIN is a more recent variant of the YIN algorithm. It was invoked in the form of a Vamp plugin[5] from within the Sonic Visualiser[6] environment, which offers the possibility to export the transcription to a MIDI file. The main difference between YIN and PYIN is that the latter has a prior for the threshold instead of a fixed value. We chose a mean of 0.15 for the beta prior, as suggested by the authors.

B. Corpus

As our goal was to develop a transcriber that captures the subtle nuances of a jazz guitar in the first place, we have created a new data corpus consisting of ten jazz guitar excerpts (phrases) from solos that were performed and recorded at Sony CSL in Paris [33]. The excerpts were annotated manually, so as to have a reference transcription. The note onsets were determined according to their perceptual salience. The same applies to the duration of single notes. The data contains both monophonic and polyphonic (chordal) melody lines, although the number of chordal passages is small. The manual transcription, which is arguably not perfect, serves as the ground truth in our experiment. The data corpus, the manual annotation, and the corresponding output from the algorithms mentioned above are all part of the electronic appendix to the manuscript. The reference transcriber and our own transcriber were tuned by hand to perform well over the entire data set.

C. Metrics

The algorithms were compared in terms of the following performance metrics. Prior to computing the scores, the provided MIDI files were converted to binary masks. The time axis was quantized.

[5]https://code.soundsoftware.ac.uk/projects/pyin
[6]http://www.sonicvisualiser.org/

F-measure In binary classification, the F-measure indicates the accuracy of a system under test. It is defined as the harmonic mean of precision and recall:

$$F = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}} \tag{43}$$

with

$$\text{precision} = \frac{|\{\text{detected pitches}\} \cap \{\text{reference pitches}\}|}{|\{\text{detected pitches}\}|} \tag{44}$$

and

$$\text{recall} = \frac{|\{\text{detected pitches}\} \cap \{\text{reference pitches}\}|}{|\{\text{reference pitches}\}|}. \tag{45}$$

Precision is the fraction of detected pitches that are in the reference transcription, i.e. which are relevant, while recall is the fraction of relevant pitches that are detected. Accordingly, precision is often interpreted as a measure of exactness or *quality* as opposed to recall, which is seen as a measure of completeness of *quantity*. In our case, a detected pitch is a point with a value of 1 in the time-pitch plane, which is equivalent to a piano-roll representation. Undetected pitches then again are points with a value of 0. The F-measure attains its best value at one and its worst value at zero. To account for the relevance of the detected pitches, we also weight each detected pitch by the MIDI velocity of the corresponding note during the computation of (44), where the velocity is given by (36). By doing so, we make sure that softer and more subtle notes, which are more likely to go unheard, have lesser impact on the precision.

Error score The error score is a common metric in speech-to-text transcription [34]. It is nothing else but the word error rate that has established itself as the ultimate accuracy metric in the field of automatic speech recognition. Respectively, the error score is computed as:

$$E = \frac{|\{\text{pitch substitutions}\}| + |\{\text{pitch deletions}\}|}{|\{\text{reference pitches}\}|} + \frac{|\{\text{pitch insertions}\}|}{|\{\text{reference pitches}\}|}. \tag{46}$$

Here again, we assess the accuracy of the algorithms based on their binary masks. Substitutions are pitch errors, such as octave or semitone errors, deletions are missing pitches, and insertions are false positives, i.e. pitches that should not be. The latter appear, e.g., when a note is detected before the actual onset, or when a note's actual duration is exceeded. Unlike the F-measure, the error score can be larger than one.

D. Results

The results of our experiment are summarized in Fig. 6. To assess the algorithms' accuracy on how well they manage to detect the correct pitch value as a function of time, we use for the time grid the duration of a thirty-second note as the step size. This applies to Fig. 6(a). For the onset detection,

a grid resolution equal to a sixteenth note is chosen. This is because we can be more tolerant to onset deviations in view of the fact that the onsets in the reference transcription may not always be exact. With that said, we can afford to be even more permissive with the note durations, since the decay of a note is not always easily defined. Therefore, we choose a time grid with a resolution of an eighth note in Fig. 6(c). On each box, the central mark is the median and the edges of the box are the 25th and 75th percentiles. The whiskers extend to the most extreme data points. The outliers are plotted individually as circles.

From Fig. 6(a) it can be seen that the proposed transcriber outperforms the reference transcriber, which performs worse[7] than PYIN, and also PYIN in terms of pitch detection ability. The pitch tracker stabilizes and further improves the accuracy of the pitchgram by some few percent: compared to Fig. 5, the box of *C* is narrower and the median moved up the scale (from 0.78 to 0.80). This is the case for both variants of the pitchgram, although the frequency-domain pitchgram yields somewhat better results. This is explained by the following.

- The time-pitch distribution is weighted *prior to* bident filtering in the case of *B*. As a consequence, we observe less octave ambiguities.
- The *implicit* detection of transients works way better for the frequency-domain pitchgram. Consequently, tremolos may not be detected equally well.
- The transcriber's empirical thresholds were tuned to the frequency-domain pitchgram in the first place. The two representations are equivalent but not identical.

The improvement from *A* to *B* is around 34 % in terms of the F-measure and it is 52 % w.r.t. the error score. Respectively, the improvements from *D* to *A* are 19 % and 45 %. Figs. 6(b) and 6(c) also show an improvement. This means that the note onset and the duration can be detected more reliably using the pitchgram. The exact performance figures for each algorithm show that the outliers stem from the third excerpt that contains a progression of chords. We would like to highlight that both the reference transcriber that uses YIN and PYIN fail almost completely on that test item, while our pitchgram transcriber achieves an F-score of 0.45. Two-thirds of the notes were left undetected, however.

Finally, to shed some light on the behavior of each tested system at the transition points between consecutive notes and to better understand their weak spots, we can analyze what the error score is actually composed of. This is depicted in Fig. 7. It can be seen that the main source of error for the reference transcriber are deletions. In part, this is an indicator of an imprecise onset detector. Also, the number of substitutions (pitch errors) near note onsets is significantly higher. PYIN, or rather its pitch tracker, tries to detect the dominant pitch, and hence it suffers mainly from deletions. Our system, then again, shows more insertions. I.e. it has detected more notes than there are according to the ground truth. These extra notes reflect string vibrations triggered by finger movements or other mechanical phenomena that make up the nuances of a live performance. For this, one could also argue that our system is better suited for live performance analysis.

[7]This is partly due to the more sophisticated pitch tracking in PYIN.

TABLE I
REAL-TIME FACTOR OF THE PROPOSED TRANSCRIPTION SYSTEM

	B	C
Frequency-domain pitchgram	0.22	—
Time-domain pitchgram	—	3.02
Incl. transcriber	0.25	3.05

The runtime[8] of our two systems *B* and *C* relative to the real time, i.e. the total signal duration, is given in Table I. It can be seen that in *B* the pitchgram computation requires less than one tenth of the time to compute the regular pitchgram. The complexity of the transcriber is negligible. Moreover, *B* runs four times faster than real time.

[8]Measured on an Intel Core i7-4510U CPU operating at 2.00 GHz.

VI. CONCLUSION

The filter bank developed in this manuscript enables a pitch-synchronous analysis of music signals. The output of the filter bank, the pitchgram, is equivalent to the real Fourier transform of a pre- and post-filtered ACF of a (wide-sense) periodic signal. Therefore, the behavior of the filter bank can be analyzed based on the classical theory on linear time-invariant systems. What is more, it renders possible the design and analysis of new and eventually more performant prototype filters. The bident is an example of a filter that can help alleviate the pitch ambiguity between a harmonic and its overtones. The proposed pitchgram may not only be useful for transcription purposes. Eventually, it may also provide the basis for statements on the expressiveness and richness of a performance, and its technical accuracy alike, see [35].

An automatic, decision-based transcription system for the pitchgram was also developed. It was validated that the system is capable of converting a jazz guitar solo into a piano roll very accurately, while taking certain subtleties (such as a high note rate, fretting noise, and pitch modulations) of a (live) performance into account. Therefore, apart from the visual examination of the pitchgram, the obtained transcription could already help carry out a basic yet robust performance analysis by means of a learning agent. Moreover, it should also be technically feasible to detect various playing techniques, which is yet to be shown. This is an outlook on future work.

The transcriber operates on an ad-hoc basis, i.e. without previous training. If combined with frequency-domain pitchgram, the transcriber has the potential to be run on-line and in real time on a low-power DSP. The faster pitchgram variant does not require an additional transient detector to identify tremolo or legato. In this all-in-one solution, everything is derived from the pitchgram. On the other hand, the time-domain pitchgram can provide cleaner tonality estimates at low frequencies in particular, with less frequency spread and blur. This is achieved with a bigger computational effort and with additional memory consumption.

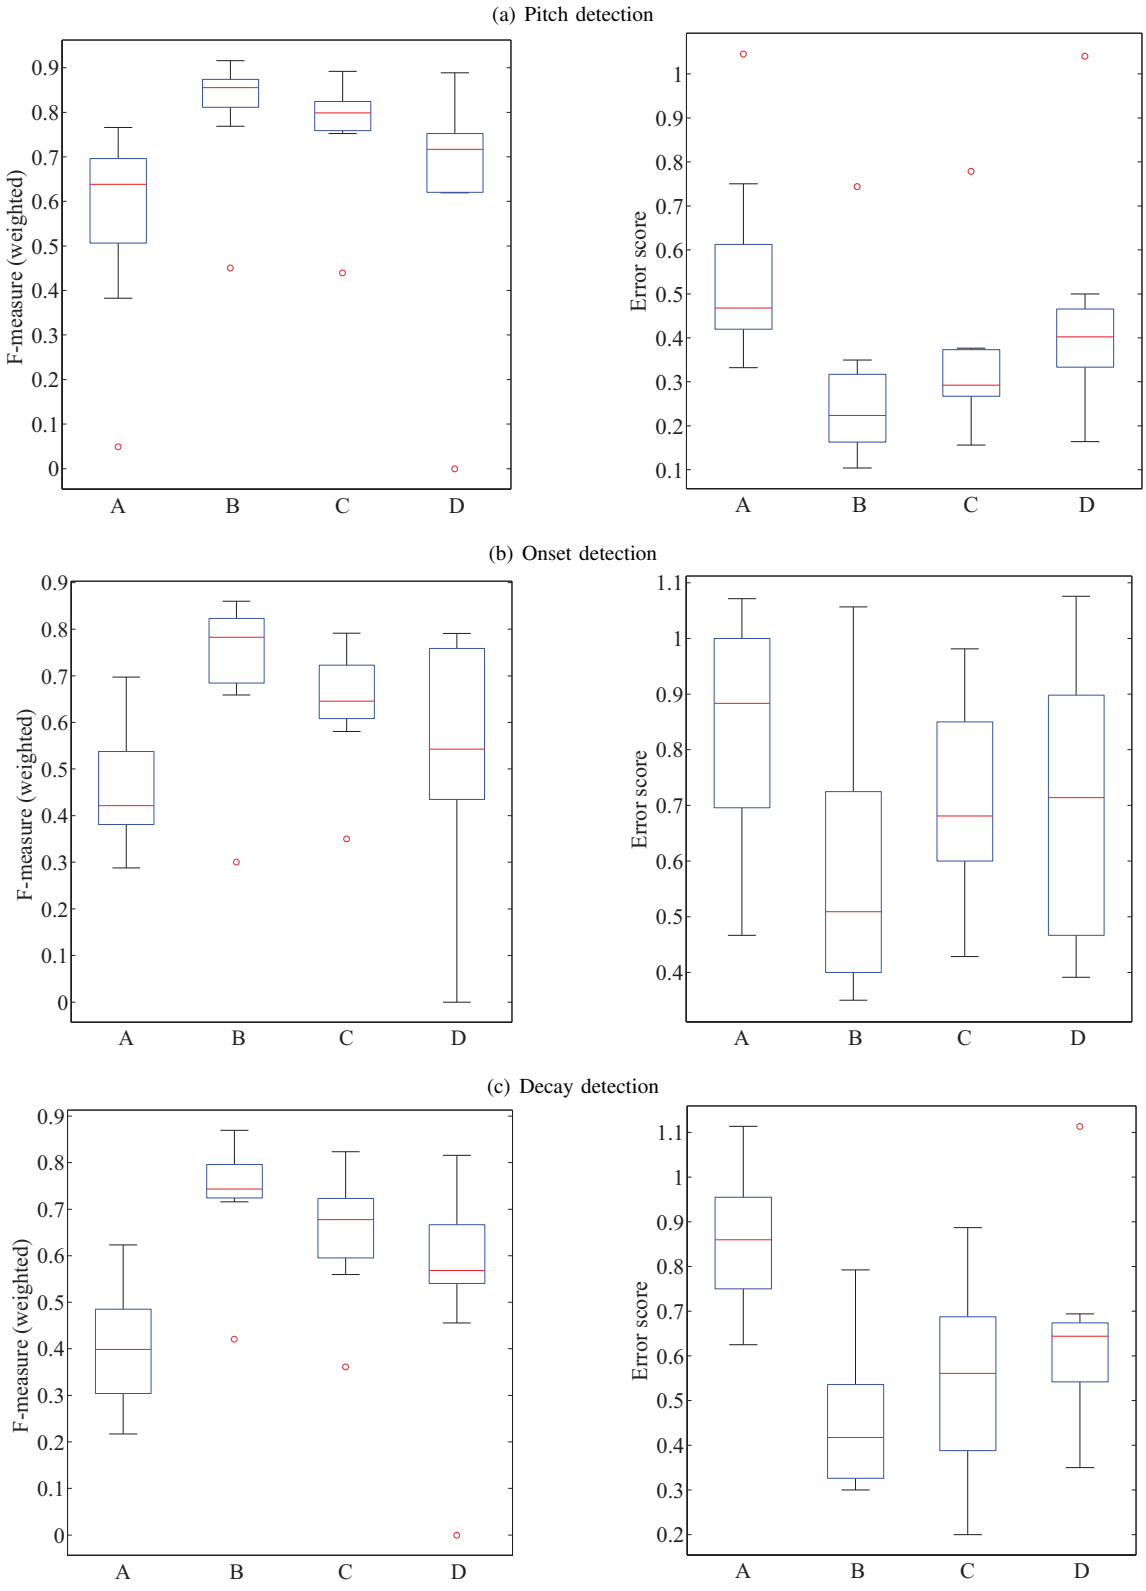

Fig. 6. Comparison of the systems under test w.r.t. the F-measure and the error score.

As a last point, we would like to underline that the major gain of our system stems from the signal representation and not so much from the pitch tracker. As the evaluation shows, with a robust signal representation and a rather simple pitch tracker, we were able to achieve a better result than the YIN-based systems, despite a transient detector or a hidden Markov model with Viterbi decoding in the post-processing stage.

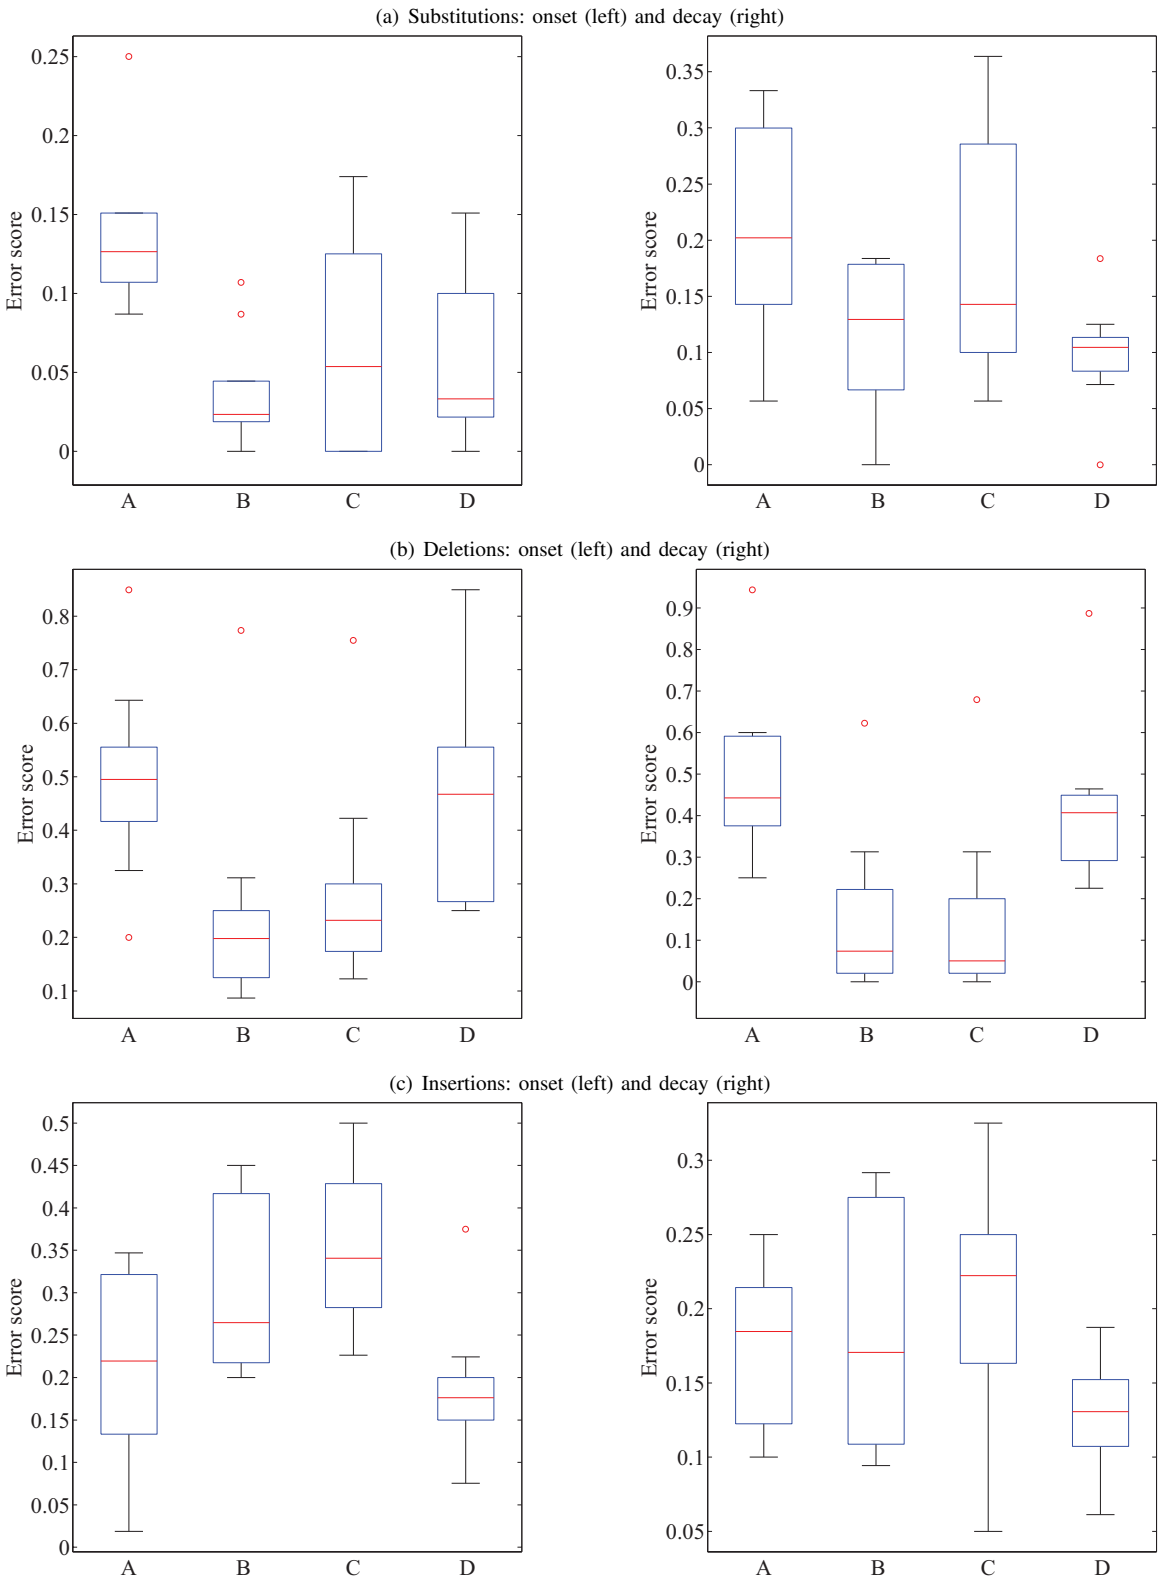

Fig. 7. Composition of the error score at transition points between notes.

ACKNOWLEDGMENT

The jazz guitar licks were performed by François Pachet. The authors cordially thank Josep L. Arcos from the "Institut d'Investigació en Intel·ligència Artificial" (IIIA) in Bellaterra for the provided support. They also thank all the reviewers of the manuscript.

REFERENCES

[1] H. Leonard, *The Real Book*, 6th ed. Hal Leonard Corp., 2005.

[2] D. J. Hermes, "Measurement of pitch by subharmonic summation," *J. Acoust. Soc. Am.*, vol. 83, no. 1, pp. 257–264, Jan. 1988.

[3] M. Goto, "A real-time music-scene-description system: predominant-F0 estimation for detecting melody and bass lines in real-world audio signals," *Speech Commun.*, vol. 43, no. 4, pp. 311–329, Sept. 2004.

[4] A. Klapuri, "Multipitch analysis of polyphonic music and speech signals using an auditory model," *IEEE Audio, Speech, Language Process.*, vol. 16, no. 2, pp. 255–266, Feb. 2008.

[5] C. Yeh, A. Roebel, and X. Rodet, "Multiple fundamental frequency estimation and polyphony inference of polyphonic music signals," *IEEE Trans. Audio, Speech, Language Process.*, vol. 18, no. 6, pp. 1116–1126, Aug. 2010.

[6] K. Dressler, "Pitch estimation by the pair-wise evaluation of spectral peaks," in *Proc. AES Intl. Conf. 42*, Jul. 2011, pp. 1–10.

[7] J. Wu, E. Vincent, S. A. Raczynski, T. Nishimoto, N. Ono, and S. Sagayama, "Polyphonic pitch estimation and instrument identification by joint modeling of sustained and attack sounds," *IEEE J. Sel. Topics Signal Process.*, vol. 5, no. 6, pp. 1124–1132, Oct. 2011.

[8] A. de Cheveigné and H. Kawahara, "YIN, a fundamental frequency estimator for speech and music," *J. Acoust. Soc. Am.*, vol. 111, no. 4, pp. 1917–1930, Apr. 2002.

[9] P. Smaragdis and J. C. Brown, "Non-negative matrix factorization for polyphonic music transcription," in *Proc. IEEE WASPAA 2003*, Oct. 2003, pp. 177–180.

[10] T. Berg-Kirkpatrick, J. Andreas, and D. Klein, "Unsupervised transcription of piano music," in *Proc. NIPS 2014*, Dec. 2014, pp. 1538–1546.

[11] B. Fuentes, R. Badeau, and G. Richard, "Harmonic adaptive latent component analysis of audio and application to music transcription," *IEEE Audio, Speech, Language Process.*, vol. 21, no. 9, pp. 1854–1866, Sept. 2013.

[12] E. Vincent, N. Bertin, and R. Badeau, "Adaptive harmonic spectral decomposition for multiple pitch estimation," *IEEE Audio, Speech, Language Process.*, vol. 18, no. 3, pp. 528–537, Mar. 2010.

[13] J. Salamon, E. Gómez, D. P. W. Ellis, and G. Richard, "Melody extraction from polyphonic music signals: Approaches, applications, and challenges," *IEEE Signal Process. Mag.*, vol. 31, no. 2, pp. 118–134, Mar. 2014.

[14] S. A. Zahorian and H. Hu, "A spectral/temporal method for robust fundamental frequency tracking," *J. Acoust. Soc. Am.*, vol. 123, no. 6, pp. 4559–4571, Jun. 2008.

[15] M. Mauch and S. Dixon, "PYIN: A fundamental frequency estimator using probabilistic threshold distributions," in *Proc. IEEE ICASSP 2014*, May 2014, pp. 659–663.

[16] C. Kehling, J. Abeßer, C. Dittmar, and G. Schuller, "Automatic tablature transcription of electric guitar recordings by estimation of score- and instrument-related parameters," in *Proc. DAFx-14*, Sept. 2014, pp. 1–8.

[17] E. Terhardt, G. Stoll, and M. Seewann, "Algorithm for extraction of pitch and pitch salience from complex tonal signals," *J. Acoust. Soc. Am.*, vol. 71, no. 3, pp. 679–688, 1982.

[18] A. H. Benade, *Fundamentals of Musical Acoustics*, 2nd ed., ser. Dover Books on Music. Dover Publications, 1990.

[19] J. Meyer and U. Hansen, *Acoustics and the Performance of Music: Manual for Acousticians, Audio Engineers, Musicians, Architects and Musical Instrument Makers*, 5th ed. Springer Science & Business Media, 2009.

[20] B. Boashash, *Time Frequency Signal Analysis and Processing: A Comprehensive Reference*, 1st ed. Elsevier, 2003.

[21] S. Saito, H. Kameoka, K. Takahashi, T. Nishimoto, and S. Sagayama, "Specmurt analysis of polyphonic music signals," *IEEE Trans. Audio, Speech, Language Process.*, vol. 16, no. 3, pp. 639–650, Mar. 2008.

[22] N. H. Fletcher, "The nonlinear physics of musical instruments," *Rep. Prog. Phys.*, vol. 62, no. 5, pp. 723–764, 1999.

[23] C. J. Plack, A. J. Oxenham, R. R. Fay, and A. N. Popper, *Pitch: Neural Coding and Perception*, ser. Springer Handbook of Auditory Research. Springer Science & Business Media, 2006, vol. 24.

[24] J. M. Barbour, *Tuning and Temperament: A Historical Survey*, ser. Dover Books on Music Series. Dover Publications, 2004.

[25] J. O. Smith, III, *Physical Audio Signal Processing*. W3K Publishing, 2010, https://ccrma.stanford.edu/~jos/pasp/.

[26] S. Gorlow and J. Janer, "Evaluation of the convolutional NMF for supervised polyphonic music transcription and note isolation," in *Proc. LVA/ICA 2015*, Aug. 2015, pp. 1–8.

[27] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, 3rd ed. Prentice Hall, 2009.

[28] J.-S. Liénard, F. Signol, and C. Barras, "Speech fundamental frequency estimation using the alternate comb," in *Proc. Interspeech 2007*, Aug. 2007, pp. 2773–2776.

[29] G. H. Wakefield, "Mathematical representation of joint time-chroma distributions," in *Advanced Signal Processing Algorithms, Architectures, and Implementations IX*, vol. 3807, Nov. 1999, pp. 637–645.

[30] T. H. Özaslan, E. Guaus, E. Palacios, and J. L. Arcos, "Identifying attack articulations in classical guitar," in *Exploring Music Contents*, ser. Lecture Notes in Computer Science, S. Ystad, M. Aramaki, R. Kronland-Martinet, and K. Jensen, Eds. Springer Berlin Heidelberg, 2011, vol. 6684, pp. 219–241.

[31] J. Serrà, T. H. Özaslan, and J. L. Arcos, "Note onset deviations as musical piece signatures," *PLoS ONE*, vol. 8, no. 7, pp. 1–9, Jul. 2013.

[32] C. Duxbury, J. P. Bello, M. Davies, and M. Sandler, "Complex domain onset detection for musical signals," Sept. 2003, pp. 1–4.

[33] S. Gorlow, M. Ramona, and F. Pachet. (2016, Nov.) A small dataset of jazz guitar licks for automatic transcription experiments. Zenodo. [Online]. Available: https://doi.org/10.5281/zenodo.167567

[34] NIST, *The 2009 (RT-09) Rich Transcription Meeting Recognition Evaluation Plan*, 2nd ed., National Institute of Standards and Technology, Feb. 2009.

[35] L. Jure, E. López, M. Rocamora, P. Cancela, H. Sponton, and I. Irigaray, "Pitch content visualization tools for music performance analysis," in *Proc. ISMIR*, Oct. 2012, pp. 1–6.